

This Page Is Inserted by IFW Operations  
and is not a part of the Official Record

## BEST AVAILABLE IMAGES

Defective images within this document are accurate representations of the original documents submitted by the applicant.

Defects in the images may include (but are not limited to):

- BLACK BORDERS
- TEXT CUT OFF AT TOP, BOTTOM OR SIDES
- FADED TEXT
- ILLEGIBLE TEXT
- SKEWED/SLANTED IMAGES
- COLORED PHOTOS
- BLACK OR VERY BLACK AND WHITE DARK PHOTOS
- GRAY SCALE DOCUMENTS

## IMAGES ARE BEST AVAILABLE COPY.

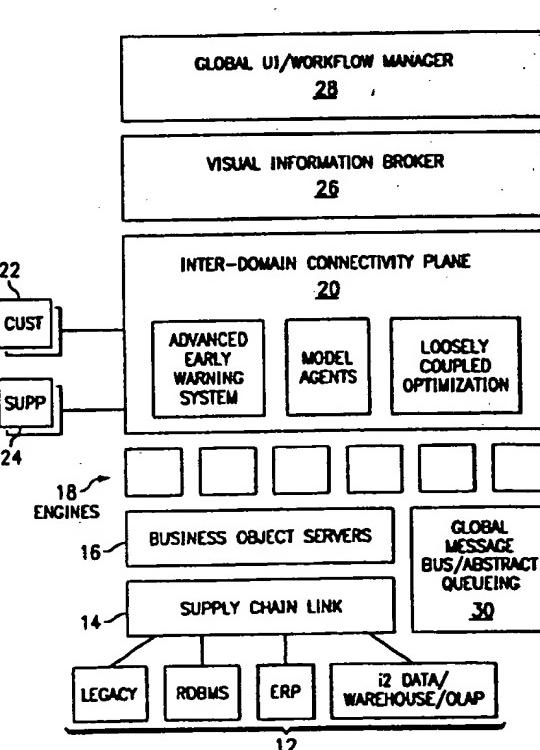
**As rescanning documents *will not* correct images,  
please do not report the images to the  
Image Problems Mailbox.**

**THIS PAGE BLANK (USPTO)**



## INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification <sup>6</sup> :  G06F 17/60		A1	(11) International Publication Number: WO 99/10825  (43) International Publication Date: 4 March 1999 (04.03.99)
<p>(21) International Application Number: PCT/US98/17359</p> <p>(22) International Filing Date: 21 August 1998 (21.08.98)</p> <p>(30) Priority Data: 08/918,227 25 August 1997 (25.08.97) US</p> <p>(71) Applicant: i2 TECHNOLOGIES, INC. [US/US]; Suite 1600, 909 E. Las Colinas Boulevard, Irving, TX 75039 (US).</p> <p>(72) Inventors: NOTANI, Ranjit, N.; 858 Kinwest Parkway #188, Irving, TX 75063 (US). MAYER, John, E.; 16323 Amberwood, Dallas, TX 75248 (US). HILERIO, Israel; 8636 Woodcreek Drive, Irving, TX 75063 (US).</p> <p>(74) Agent: MILLS, Jerry, W.; Baker &amp; Botts, L.L.P., 2001 Ross Avenue, Dallas, TX 75201-2980 (US).</p>		<p>(81) Designated States: AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, CA, CH, CN, CU, CZ, DE, DK, EE, ES, FI, GB, GE, GH, GM, HR, HU, ID, IL, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MD, MG, MK, MN, MW, MX, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TR, TT, UA, UG, UZ, VN, YU, ZW, ARIPO patent (GH, GM, KE, LS, MW, SD, SZ, UG, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GW, ML, MR, NE, SN, TD, TG).</p> <p><b>Published</b> With international search report. Before the expiration of the time limit for amending the claims and to be republished in the event of the receipt of amendments.</p>	
<p>(54) Title: SYSTEM AND PROCESS FOR INTER-DOMAIN INTERACTION ACROSS AN INTER-DOMAIN CONNECTIVITY PLANE</p> <p><b>(57) Abstract</b></p> <p>A computer implemented system and process enables inter-domain interaction. The system includes a native plane (292) that has a first native source (294) associated with a first domain and a second native source (296) associated with a second domain. The first and second native sources (294, 296) include domain related data and applications. The system also includes an inter-domain connectivity plane (298) that has a first domain engine (300) associated with the first domain, and a second domain engine (302) associated with the second domain. The first and second domain engines (300, 302) operate to perform domain analysis functions. The system further includes adapters associated with the first and second native sources (294, 296) where the adapters operate to abstract data and information from the first and second native sources (294, 296) onto the inter-domain connectivity plane (298). The first and second domain engines (300, 302) respectively receive data and information abstracted from the first and second native sources (294, 296). The first and second domain engines (300, 302) also operate to perform global messaging (306) between one another transparently supported by native messaging (308) on the native plane (292).</p>			



**FOR THE PURPOSES OF INFORMATION ONLY**

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AL	Albania	ES	Spain	LS	Lesotho	SI	Slovenia
AM	Armenia	FI	Finland	LT	Lithuania	SK	Slovakia
AT	Austria	FR	France	LU	Luxembourg	SN	Senegal
AU	Australia	GA	Gabon	LV	Latvia	SZ	Swaziland
AZ	Azerbaijan	GB	United Kingdom	MC	Monaco	TD	Chad
BA	Bosnia and Herzegovina	GE	Georgia	MD	Republic of Moldova	TG	Togo
BB	Barbados	GH	Ghana	MG	Madagascar	TJ	Tajikistan
BE	Belgium	GN	Guinea	MK	The former Yugoslav Republic of Macedonia	TM	Turkmenistan
BF	Burkina Faso	GR	Greece	ML	Mali	TR	Turkey
BG	Bulgaria	HU	Hungary	MN	Mongolia	TT	Trinidad and Tobago
BJ	Benin	IE	Ireland	MR	Mauritania	UA	Ukraine
BR	Brazil	IL	Israel	MW	Malawi	UG	Uganda
BY	Belarus	IS	Iceland	MX	Mexico	US	United States of America
CA	Canada	IT	Italy	NE	Niger	UZ	Uzbekistan
CF	Central African Republic	JP	Japan	NL	Netherlands	VN	Viet Nam
CG	Congo	KE	Kenya	NO	Norway	YU	Yugoslavia
CH	Switzerland	KG	Kyrgyzstan	NZ	New Zealand	ZW	Zimbabwe
CI	Côte d'Ivoire	KP	Democratic People's Republic of Korea	PL	Poland		
CM	Cameroon	KR	Republic of Korea	PT	Portugal		
CN	China	KZ	Kazakhstan	RO	Romania		
CU	Cuba	LC	Saint Lucia	RU	Russian Federation		
CZ	Czech Republic	LI	Liechtenstein	SD	Sudan		
DE	Germany	LK	Sri Lanka	SE	Sweden		
DK	Denmark	LR	Liberia	SG	Singapore		
EE	Estonia						

SYSTEM AND PROCESS FOR INTER-DOMAIN INTERACTION  
ACROSS AN INTER-DOMAIN CONNECTIVITY PLANE

TECHNICAL FIELD OF THE INVENTION

This invention relates in general to the field of supply chain, enterprise and site planning, and more particularly to a system and process for inter-domain interaction, processing and communication across an inter-domain connectivity plane.

BACKGROUND OF THE INVENTION

Supply chain, enterprise and site planning applications and environments are widely used by manufacturing entities for decision support and to help manage operations. Decision support environments for supply chain, enterprise, and site planning have evolved from single-domain, monolithic environments to multi-domain, monolithic environments. Conventional planning software applications are available in a wide range of products offered by various companies. These decision support tools allow entities to more efficiently manage complex manufacturing operations. However, supply chains are generally characterized by multiple, distributed and heterogeneous planning environments. Thus, there are limits to the effectiveness of conventional environments when applied to the problem of supply chain planning due to monolithic application architectures. Further, these problems are exacerbated when there is no one "owner" of the entire supply chain.

It is desirable for the next evolutionary step for

planning environments to establish a multi-domain, heterogenous architecture that supports products spanning multiple domains, as well as spanning multiple engines and products. The integration of the various planning environments into a seamless solution can enable inter-domain and inter-enterprise supply chain planning. Further, an important function provided by some planning applications is the optimization of the subject environment rather than simply tracking transactions. In particular, the RHYTHM family of products available from I2 TECHNOLOGIES provide optimization functionality. However, with respect to planning at the enterprise or supply chain level, many conventional applications, such as those available from SAP, use enterprise resource planning (ERP) engines and do not provide optimization. It is thus also desirable to expand planning analysis and optimization to the inter-enterprise or inter-domain level to enable planning optimization across the supply chain.

20 SUMMARY OF THE INVENTION

In accordance with the present invention, a system and process for inter-domain interaction across an inter-domain connectivity plane are disclosed that provide advantages over conventional supply chain, enterprise and site planning environments.

According to one aspect of the present invention, a computer implemented system enables inter-domain interaction. The system includes a native plane that has a first native source associated with a first domain and a second native source associated with a second domain. The first and second native sources include domain related data and applications. The system also includes an inter-domain

connectivity plane that has a first domain engine associated with the first domain and a second domain engine associated with the second domain. The first and second domain engines operate to perform domain analysis functions. The system further includes adapters associated with the first and second native sources where the adapters operate to abstract data and information from the first and second native sources onto the inter-domain connectivity plane. The first and second domain engines respectively receive data and information abstracted from the first and second native sources. The first and second domain engines also operate to perform global messaging between one another transparently supported by native messaging on the native plane.

According to another aspect of the present invention, a computer implemented process is provided for inter-domain interaction. A native plain is established that has a plurality of native sources respectively associated with domains where the plurality of native sources include domain related data and applications. An inter-domain connectivity plane is also established that has a plurality of domain engines respectively associated with the domains where the plurality of domain engines operating to perform domain analysis functions. Data and information is abstracted from the plurality of native sources onto the inter-domain connectivity plane and is received at the plurality of domain engines. Further, messaging is performed between the plurality of domain engines transparently supported by native messaging on the native plane.

A technical advantage of the present invention is that the inter-domain connectivity plane provides an abstracted

layer to allow harmonization across distributed domains and provides significant advantages by harmonizing such aspects as object structure, messaging paradigm, naming and security.

5 Additional technical advantages should be readily apparent to one skilled in the art from the following figures, descriptions, and claims.

BRIEF DESCRIPTION OF THE DRAWINGS

10 A more complete understanding of the present invention and advantages thereof may be acquired by referring to the following description taken in conjunction with the accompanying drawings, in which like reference numbers indicate like features, and wherein:

15 FIGURE 1 is a block diagram providing an overview of one embodiment of a core environment that enables supply chain analysis and optimization according to the present invention;

20 FIGURES 2A and 2B are block diagrams illustrating conventional many-to-many conversion for inter-domain analysis and one-to-many conversion for inter-domain analysis according to the present invention;

25 FIGURE 3 is a block diagram of one embodiment of relationships among adaptation dimensions associated with inter-domain analysis and optimization according to the present invention;

FIGURE 4 is a block diagram of the visual information broker (VIB) operating as a middle tier to various engines and other data sources according to the present invention;

30 FIGURE 5 is a block diagram showing the VIB interaction with various data sources as well as browser and non-browser user interfaces according to the present

invention;

FIGURE 6 is a block diagram of a business object server operating as data server according to the present invention;

5 FIGURE 7 is a block diagram showing different components of queuing and the global messaging bus according to the present invention;

10 FIGURE 8 is a block diagram showing transactional messaging between applications according to the present invention;

FIGURES 9A and 9B are block diagrams of a naive approach to inter-domain analysis and optimization and of the use of model agents as partial replicas according to the present invention;

15 FIGURES 10A and 10B are block diagrams of many-to-many interaction between domains and of simplified domain interaction enabled by the inter-domain connectivity plane according to the present invention;

20 FIGURE 11 is a block diagram of one embodiment of an inter-domain connectivity plane according to the present invention; and

25 FIGURE 12 is a block diagram of a hub and spoke collaboration network formed by engines on the inter-domain connectivity plane.

#### DETAILED DESCRIPTION OF THE INVENTION

FIGURE 1 is a block diagram providing an overview of one embodiment of a core environment, indicated generally at 10, that enables supply chain analysis and optimization according to the present invention. Core environment 10 provides a framework for inter-domain and inter-enterprise analysis and optimization. This framework allows

optimization in the extended enterprise distributed environment. The framework also provides an inter-domain decision support environment that allows optimization for the extended enterprise. The framework further allows the embedding of other applications within it and can be embedded within other applications. Core environment 10 integrates multiple engines along multiple dimensions including user interface (UI), messaging, data access, data modeling and other dimensions. Core environment 10 provides scalability, common protocols and security for supply chain planning and provides a universal adapter framework to provide connectivity between components and with other applications. All of the components and functionality within core environment 10, which are discussed herein, generally can be implemented by software executing on a computer device comprising typical components such as a processor, memory, input/output (I/O) devices and data storage devices.

As shown, core environment 10 includes a number of data sources 12 associated with planning environments that can be used by various enterprises and sites. Data sources 12 include, but are not limited to, legacy data sources (e.g., IBM mainframe and mid-range systems), relational database management systems (RDBMS) (e.g., ORACLE, SYBASE), enterprise resource planning systems (ERP) (e.g., SAP), I2 TECHNOLOGIES data sources, data warehouses and on-line analytical processing (OLAP) sources. Core environment 10 includes a supply chain data link 14 that provides a layer to interface to data sources 12. Supply chain data link 14 can, for example, be a layer established by the RHYTHM-LINK product available from I2 TECHNOLOGIES.

Business object servers (BOS's) 16 work through supply

chain link 14 to interface with associated data sources 12. Alternatively, BOS's 16 can interface directly to data sources 12. For each data source 12, BOS's 16 use a dynamically loaded adapter to interface with the particular 5 data source 12 and use a common BOS interface to communicate to higher levels within core environment 10. The structure and operation of BOS's 16 are also shown and described below. As shown in FIGURE 1, BOS's 16 can interface to planning engines 18 which comprise various 10 domain engines that handle planning analysis and optimization across the supply chain.

Planning engines 18 are integrated on an inter-domain connectivity plane 20. Inter-domain connectivity plane 20 provides an abstract layer for messaging and transfer of 15 model agents that allow engines 18 to operate at the supply chain level. Inter-domain connectivity plane 20 is specifically geared towards inter-domain decision support, and the structure and operation of inter-domain connectivity plane 20 and its components are also described 20 below. Among the functions provided by inter-domain connectivity plane 20 are messaging between domains, an advanced early warning system between domains for emergency events, the transfer of model agents and the enablement of loosely coupled optimization. Inter-domain connectivity 25 plane 20 also allows interaction with custom applications 22 and support applications 24.

A visual information broker (VIB) 26 interfaces to inter-domain connectivity plane 20 to obtain information from various sources and package that information into a 30 common format. The common format allows global user interface (UI) 28 to use a library of interface components that does not need to be changed every time the information

source changes. In one embodiment, user interface components are implemented using JAVABEANS or ACTIVEDEX user interface components, and the JAVA language is used to provide extensibility. Analogous to BOS's 16, visual information broker 26 uses dynamically loaded adapters specifically designed to interface to particular sources of information. Core environment 10 also includes a global message bus and abstract queuing 30 which enables interaction between engines 18 on inter-domain connectivity plane 20. In one implementation, synchronous communications are handled by common object request broker architecture (CORBA) and distributed component object model (DCOM), and asynchronous communications are handled by supporting many messaging layers, including, but not limited to, TIBCO, MQSERIES, ACTIVEWEB and SMTP (simple mail transfer protocol).

Within core environment 10, global user interface 28 supports a common user interface across multiple engines 18 and sources 12. Global user interface 28 allows a common set of reusable, data-aware components to be established across the multiple engine types. A number of data-aware user interface components can be supported such as: tree, tree-table, table, two-dimensional and three-dimensional charts (line, bar, Gantt, pie), maps and graphs, multi-dimensional, and axis-cross. Global user interface 28 also can establish a common extensibility mechanism, a common security framework and a common user model.

Core environment 10 can support a number of distinct kinds of user interfaces, including, for example, an ACTIVEDEX user interface, a JAVABEANS-based stand alone user interface, a JAVABEANS-based browser user interface and a dynamic-SHTML servelet user interface. The ACTIVEDEX user

interface can be supported via a COM interface to a data manager, which can be a sub-component of visual information broker (VIB) 26. Also, data aware ACTIVEX components that are aware of VIB 26 data models can be provided. This user 5 interface would also support JAVABEANS components wrapped as ACTIVEX. The JAVABEANS-based stand alone user interface can include JAVABEANS that are aware of VIB 26 data models. This user interface would support ACTIVEX components wrapped as JAVA components using the MICROSOFT JAVA 10 virtual machine. The JAVABEANS-based browser user interface can be a pure JAVA user interface. It can be designed to run in any browser as well as on network computers. As a pure JAVA interface, the user interface will run inside any JAVA virtual machine (JVM). The 15 dynamic-SHTML servelet user interface can be designed for low-bandwidth situations such as those found on the Internet. This approach distributes the processing at the server side producing dynamic HTML pages that are displayed on a web browser.

20 In general, core environment 10 is designed to integrate existing and future planning engines at multiple levels and possesses a number of key attributes. Core environment 10 is designed to be open and establish protocols by which applications can be developed to operate 25 under core environment 10. Core environment 10 is component based and uses standard component architectures. For example, CORBA and DCOM can be used as a distributed component architecture, and JAVABEANS and ACTIVEX can be used as user interface component architecture. By having 30 a component architecture, core environment 10 allows changes to be released in pieces and not all at once, allows users to select and use only components that are

useful, allows components to be re-used, and allows third party components to be incorporated.

Core environment 10 can be scaleable in a variety of different ways. Components can be load balanced and multi-threaded which leads to higher throughput. As described 5 below, the universal adaptor framework avoids the many-to-many issue which leads to greater scalability. Further, security is handled by core environment 10 so that individual engines do not have to be worried about 10 security. This also leads to greater scalability.

With respect to security, core environment 10 can provide a number of forms of security. For example, there can be security between user interface 28 and visual information broker 26 and messaging security. In both 15 cases, both encryption as well as user authentication can be provided. An important aspect is that core environment 10 can implement security in such a way that it does not require each and every engine to be concerned with security.

Core environment 10 may, however, place some minimal requirements on the engine architecture. In one implementation, core environment 10 makes the assumption that an interface can be provided to the engine that is either an external interface or an in-memory JAVA 25 interface. In this implementation, if an external interface is provided, the interface can be a CORBA interface, a COM/DCOM interface, a shared library, structured query language (SQL), an object database interface or a socket interface. If an in-memory JAVA interface is provided, the interface can typically be accomplished by embedding a JAVA virtual machine (JVM) in 30 the engine and writing a local JAVA interface (LJI) that

calls the local native interface (LNI) via the JAVA native interface (JNI). Another assumption that can be made about the engines is that the engines have an ability to be a CORBA client and have the ability to marshall CORBA structures. Native components can be provided for both of these tasks. However, in this case, core environment 10 specifically does not make the requirement that an engine must be a CORBA server (which would require hooking up the engine's event loop with the CORBA event loop).

10

#### Universal Adapter Framework

One problem in creating core environment 10 is matching an interface (along some dimension) of one product or application with that of another. One possible solution 15 is to write converters between all possible combinations of interfaces. However this leads to a classic many-to-many problem where the number of converters is unmanageable. FIGURE 2A is a block diagram illustrating this conventional many-to-many conversion problem for inter-domain analysis. As shown, an environment 40 may contain a plurality of products, applications, data sources or other entities 42 that need to interface with one another. In environment 40, there is a converter, represented by the lines, associated with the interface between each pair of entities 25 42. As can be seen, this many-to-many solution will generate a requirement to create and manage a prohibitively large number of converters as the number of entities 42 increase.

FIGURE 2B is a block diagram illustrating one-to-many conversion for inter-domain analysis according to the 30 present invention. This solution is enabled by a universal adaptor framework in which the different adaptation

dimensions use the same mechanism for adaptation. As shown, an environment 44 can contain a plurality of products, applications, data sources or other entities 46. However, unlike environment 40 of FIGURE 2A, environment 44 further includes a universal interface 50. With this scheme, only one adapter 48 needs to be created and managed between each entity 46 and universal interface 50. In general, an adaptor 48 maps a proprietary interface from an entity 46 into universal interface 50. This allows a one-to-many conversion instead of the many-to-many conversion of FIGURE 2A. Further, adapters 48 can be dynamically loadable such that an appropriate adapter 48 can be loaded when an interface to a particular entity 46 is needed and then removed if no longer needed.

This universal adaptor framework works across multiple dimensions using dynamically loaded adaptors 48. These adaptation dimensions can include visual (handled by VIB), data (handled by BOS's), messaging and queuing. Dynamically loaded adapters 48 help to establish openness within core environment 10 and can use the same mechanism across all adaptation dimensions. This reduces the learning curve for users and developers within core environment 10.

25 Adaptation Dimensions

FIGURE 3 is a block diagram of one embodiment of relationships among adaptation dimensions associated with inter-domain analysis and optimization according to the present invention. FIGURE 3 shows an example environment, indicated generally at 60, that includes four adaptation dimensions -- user interface, data, queuing and message bus -- all of which involve dynamically loaded adapters.

Beginning with data sources, environment 60 includes a planning engine 62 and data storage 64 accessed by a business object server (BOS) 66 using dynamically loadable adapters 68 and 70. An SAP based system 72 and an ORACLE based system 74 are linked through a RHYTHM-LINK application 76. A second BOS 78 then accesses both data storage 64 and RHYTHM-LINK application 76 using adapters 80 and 82. Similarly, a common data model storage 84 is accessed by a BOS 86 and a BOS 90 using adapters 88. BOS's 66, 78, 86 and 90 all interface to a domain engine 92, which can be one of a number of domain engines operating within a supply chain. In addition, RHYTHM-LINK application 76 can be enabled to interface directly to domain engine 92, and in particular can do so when domain engine 92 is an I2 TECHNOLOGIES' product. Domain engine 92, in turn, provides information to a visual information broker 94 through dynamically loadable adapters 96. Visual information broker then provides data to a common user interface (UI) 98 for providing a display to a user.

Environment 60 also includes a global message bus 100 which uses adapters 102 to interface to native messaging functionality and allow engine 92 to interact with other domain engines. Global message bus 100 also interfaces with a queue adapter 104 that uses adapters 106 to allow messages to be persisted on either a relational database management system (RDBMS) 108 or an object oriented database management system (ODBMS) 108. This persistence can form the basis for guaranteed message delivery.

The user interface adaptation dimension concerns the presentation of information to users and involves common user interface 98, visual information broker 94 and adapters 96. For example, if an engine 92 needs to be

accessible from common user interface 98, then the engine 92 should have visual information broker (VIB) adaptors 96 created for it. Likewise, although not shown, if an application needs to embed an engine 92 into the 5 application's user interface, then the application's user interface can make a call to visual information broker 94 which, in turn, interfaces to engine 92.

The data adaptation dimension concerns the accessibility of data to an engine 29 and involves BOS's 10 66, 78, 86 and 90 and BOS adapters 68, 70, 80, 82 and 88. If an application or data source needs to be accessible from engine 92, then there should be business object server (BOS) adaptor created for the application or data source.

The queue adaptation dimension concerns queues for 15 transactions and is not directly illustrated in FIGURE 3. Queue adaptors, for example, allow queues to be built on transactional databases, including relational database management systems (RDBMS's), such as ORACLE and SYBASE, and object data base management systems (ODBMS's). Queue 20 adapters can be built to any transactional database a main application or engine is using. The queue adaptation layer is also shown and described below.

The message bus adaptation layer concerns communication over message layers and involves global 25 message bus 100 and adapters 102. For example, if an application needs to have communication to an engine 92 implemented over the application's native message layer, then an adaptor 102 can be created to that native message layer.

#### Visual Information Broker

FIGURE 4 is a block diagram of the visual information

broker (VIB) operating as a middle tier to various engines and other data sources according to the present invention. As shown, a first engine 110 (of type 2) has an engine interface 112, and a second engine 114 (of type 1) has an engine interface 116. The visual information broker (VIB) 118 accesses engine interface 112 using an adapter 120 and accesses engine interface 116 using an adapter 122. Visual information broker 118 can then provide information to a user interface 124 which can include a local caching proxy server 126.

VIB 118 then has a common interface which can be accessed by user interface 124 regardless of the type of engine from which VIB 118 obtains information. As mentioned above, user interface 124 can thus include a library of data-aware components such as ACTIVEX and JAVABEANS components. VIB 118 can route incoming requests from user interface 124 to engines 110 and 114 which are of different types. VIB 118 can also load-balance engines of the same type. Such a data manager sub-component of VIB 118 is a part of the universal adaptor framework described above. VIB 118 and adapters 120 and 122 allow user interface-oriented data models from multiple engines 110 and 114 to be adapted into a common data model oriented for user interface 124. This forms a basis for common user interface 124.

VIB 118 can provide an interface across multiple sources including databases, in-memory engines, CORBA servers, flat-files, messaging, object databases, etc. VIB 118 uses dynamically loadable adapters as appropriate to interface with the sources. The adapters can be specifically designed to connect to desired source and to VIB 118. This scheme allows support of a wide variety of

data models, including tables, trees, name-value pairs, multi-dimensional, and object graphs.

FIGURE 5 is a block diagram showing interaction between the visual information broker and various data sources as well as browser and non-browser user interfaces according to the present invention. As shown, the environment can include a browser user interface 130 and a non-browser user interface 132. Browser user interface 130 connects through a workstation 134 that provides a firewall 136 and an IIOP tunnel 138.

A visual information broker (VIB) 140 can provide access to information sources 142 (DCOM, CORBA) for both IIOP tunnel 138 and non-browser user interface 132. VIB 140 can also interface with a VIB 144. Another VIB 146 can provide access to information sources 148 (SQL, OLAP, I2-FP, I2-SCP) for non-browser user interface 132. As with VIB 140, VIB 146 can also interface with VIB 144. VIB's 140, 144 and 146 can be designed to access multiple data sources simultaneously, and the data sources can include: CORBA servers, DCOM servers, COM objects, RDBMS data, ODBMS objects and in-memory objects. Further, VIB's 140, 144 and 146 can be load-balanced to divide the load across multiple servers or other sources of information, and this load-balancing can be transparent to user interface 130 or 132.

25

Visual information brokers can support a multi-threaded environment with a thread-pool model where individual requests can be executed in their own thread until a certain maximum number after which they can be queued. Visual information brokers also can be directly embedded in an engine and can run on the same machine as the engine for high-throughput compared to network access.

In one embodiment, visual information brokers include data manager, event channel and page manager sub-components. The data manager sub-component can be accessible via CORBA as well as COM and supports multiple data models including: tree-table, name-value list, axis-cross, multi-dimensional and object graph. The data manager can also support client-originated execution of model agents. The event channel sub-component allows the common user interface to subscribe to and be notified of server events. The page manager sub-component can allow a user interface to be created in an SHTML format. This format allows inter-leaving of static HTML with dynamic servelets. It is more efficient than CGI since it does not require a separate process spawned per request, but only a separate thread per request (up to a maximum since it uses a thread-pool model). The page manager also can be compatible with the JAVASOFT servelet application program interface (API). It can load in standard servelet components as well as application-specific servelets. The page manager can be similar to a JAVASOFT JAVASERVER with a number of additional features and differences. The page manager can be instantiated stand-alone or within another process. In process instantiation allows a very tight coupling between the servelet and the process. The page manager can be load-balanced and is designed to work with a web server and not be a replacement for one. The page manager also links servelet parameterization with page parameterization so that dynamic servelets can form direct parameterized links to SHTML pages.

30

Business Object Servers

FIGURE 6 is a block diagram of a business object

server (BOS) operating as a data server according to the present invention. A number of data sources 150 have associated dynamically loadable BOS adapter's 152 that interface between data sources 150 and a business object server (BOS) 154. Business object server 154 includes a BOS interface 156 which can be standard across business object servers. A BOS client 158 can then interface with business object server 154 through BOS interface 156. As shown, BOS client 158 can read from and write to data sources 150 through business object server 154, and business object server 154 can handle the specific interface with data source 150 via an appropriate BOS adapter 152. The dynamically loadable nature of BOS adapters 152 mean that they can be loaded as needed and then removed without having to otherwise affect the operation of business object server 154.

In general, BOS client 158 can be a planning engine, and business object server 154 can serve up objects to the engine from multiple different data sources 150. In this manner, business object servers 154 form an integral part of the described universal adaptor framework. BOS adapters 152 adapt data from the multiple data sources 150 into specific business objects. Each business object server can thus be classified by the data source interfaces to which it adapts.

BOS adapters 152 to data sources 150 can have a common model and an other model format. For the common model format, a standard BOS interface 156 and standard BOS adaptors 152 can be fused into a single object. Effectively, in this case, BOS adaptors 152 then provide only an object-relational mapping. For non-standard BOS adapters 152, an application can be provided to allow

development of special business object servers 154 as well as BOS adaptors 152. For the other model format, there can be BOS adaptors 152 to many non-standard data sources 150, including enterprise resource planning (ERP) systems and 5 other planning systems. BOS adaptors 152 can operate to make accessible proprietary non-standard data sources 150. For the other model format, an application can, again, be provided to allow development of special business object servers 154 and BOS adaptors 152.

10

#### Queuing and Messaging

FIGURE 7 is a block diagram showing different components of queuing and the global messaging bus according to the present invention. The queuing and global 15 messaging are also built upon the described universal adaptor framework. With respect to queuing, a storage 160 can contain queues 162 that include transactions for a relational or object database management system (RDBMS, ODBMS). A dynamically loadable queue adapter 164 provides 20 an interface from a queue manager 166 to queues 162. Queue manager 166 includes a queue application program interface (API) 168 that can be standard across queue managers 166. A transactional message manager (TMM) 170 and an application 172 can interface to queue manager 166 via 25 queue API 168.

Similarly, with respect to global messaging, transactional message manager (TMM) 170 and application 172 can interface through a message bus API 174 to a message bus manager 176. Message bus manager 176, in turn, can use 30 a dynamically loadable message bus adapter 178 to interface to native messaging 180 of the underlaying native applications. This allows global messaging to be built on

top of any third party messaging solutions including, for example, ACTIVEWEB, TIBCO, MQSERIES, NEONET, SMTP and FTP. This also allows applications 172 to be abstracted from the underlying native message layer 180. In one embodiment, 5 three levels of messaging are provided having different characteristics: fast/reliable messaging; certified messaging; and guaranteed/transactional messaging. The fast/reliable messaging can provide a reasonable efforts attempt to deliver the messages. The certified messaging 10 can provide certifications of message delivery to the client application. Third, guaranteed/transactional messaging can be provided between queues on transactional databases to ensure delivery of messages, including messaging between different transactional databases (e.g., 15 between ORACLE and SYBASE databases or between RDBMS and ODBMS databases).

FIGURE 8 is a block diagram showing transactional queuing and messaging between applications according to the present invention. As shown, a first application 190 has 20 a transactional context that can include a storage 192 holding relational database information 194 and transactional queues 196. A queue manager 198 provides an interface to queues 198 for application 190 as well as a transactional message manager (TMM) 200. Within a 25 messaging transactional context, a message bus manager 202 provides an interface between transactional message manager 200 and native messaging layer 204.

Similarly a second application 206 has a transactional 30 context that can include a storage 208 holding object database information 210 and transactional queues 212. A queue manager 214 provides an interface to queues 212 for application 206 as well as a transactional message manager.

(TMM) 216. Within the messaging transactional context, a message bus manager 218 provides an interface between transactional message manager 216 and native messaging layer 204.

5 This framework allows an abstract queuing and global messaging layer between applications 190 and 206 to be supported on an underlying native messaging layer 204. Within this messaging framework, point-to-point messaging with global addressing, publish and subscribe messaging and efficient encrypted, user-authenticated messaging can be supported. Further, an out-of-the-box solution consisting of message bus adaptors to ACTIVEWEB or TIBCO, and queue adaptors on top of an ODBC compliant database can be provided. These adaptors can be bundled along with the 10 messaging solution (ACTIVEWEB or TIBCO) to create the out-of-the-box messaging solution.

15

#### Model Agents

The core environment is an inter-domain architecture 20 that allows, for example, inter-enterprise optimization. Aspects of the environment that address a need for an inter-domain solution include: a security framework (both user interface to engine and messaging), messaging that is global in scope (including a global naming scheme, and 25 transactional messaging to allow multiple transactional domains to be kept transactionally consistent). Further, inter-domain analysis and optimization are enabled by allowing domains to exchange model agents that comprise partial replicas of the remote domains.

30 FIGURE 9A is a block diagram of a naive approach to inter-domain analysis. As shown, an environment 220 can include a plurality of domains 222, 224 and 226 (domains A,

B and C). The naive approach is characterized by a distributed analysis, represented by arrow 228, in which distributed algorithms are run directly across the multiple domains 222, 224 and 226. This approach has some severe drawbacks. For example, reliability falls off dramatically as the number of domains increases. Also, in this approach, permissibility (i.e., one domain allowing access by another domain to its model) and security need to be intrinsically bound up with the distributed analysis. This places a heavy burden on the analysis algorithm and would function too slowly but for the simplest of analysis algorithms. Another alternative, not shown, is to have every domain 222, 224 and 226 independently model all of the other domains 222, 224 and 226. This approach, however, suffers from the difficulty of accurately knowing the model and data for other domains 222, 224 and 226. This approach also suffers from a scalability problem with going to large numbers of domains 222, 224 and 226.

FIGURE 9B is a block diagram of the use of model agents as partial replicas of remote domains according to the present invention. The model agents provide an alternative mechanism that makes feasible inter-domain analysis and optimization. As shown, an environment 230 can include a plurality of domains 232, 234 and 236 (domains A, B and C). These domains 232, 234 and 236 each have models that are used for local planning analysis and optimization. According to the present invention, domains 232, 234 and 236 transfer model agents to each other that are partial replicas of the respective domains' model. The transfer of these model agents is preferably via a push scheme to subscribing domains, but may involve pull schemes or other transfer schemes. For example, domain 232 (domain

A) provides a model agent 238 to domain 236 (domain C) that is a partial replica of the model for domain 232. Model agent 238 comprises that portion of the model for domain 232 that is needed by domain 236 and to which domain 232 desires to grant external access. Similarly, domain 234 (domain B) can provide a model agent 240 to domain 236 (domain C) that is a partial replica of the model for domain 234. Again, model agent 240 comprises that portion of the model for domain 234 that is needed by domain 236 and to which domain 234 desires to grant external access. As a result, instead of running a distributed analysis over multiple domains, model agents 238 and 240 allow a local analysis, as represented by arrow 242, to be run on domain 236 using a combination of the local model (for domain C) and portions of the remote models (for domains A and B) as represented by model agents 238 and 240.

As should be understood, the use of model agents as partial replicas of remote domains allows each interested domain to accomplish inter-domain analysis and optimization via local processes. Additionally, mechanisms can be provided to continually replenish the model agents from remote domains by pushing the model agents to all subscribing domains, or otherwise providing updates to other domains. The model agents are thus partial replicas of models that can be passed around networks as discrete packages. It should be understood that the "partial" aspect can be important since typically a domain would wish to send another domain only a portion of its complete model rather than the whole thing.

In one embodiment, there are three levels of model agents, each with increasing degrees of sophistication. The first level is data model agents. These are the

simplest of the three and are used to pass around pure data. The second level of model agents are object model agents. These model agents allow entire graphs of objects to be passed around. Allowing graphs of objects to be passed around enables sophisticated forms of collaboration that would be difficult to achieve only with data model agents. For example, using object model agents, one domain can send another a partial picture of its supply chain model. The third level of model agents are behavior model agents. These model agents are more sophisticated and allow entirely new behaviors to be passed from one domain to another. For example, one domain could send another domain a new pricing policy or inventory policy as a behavior model agent.

The model agents of the present invention solve both the infeasibility and scaleability problems of other inter-domain approaches. The model agents allow a needed portion of a remote domain to reside locally and be used for analysis and optimization. The model agents are generally extracted by each domain and provided to other subscribing domains. A subscribing domain can obtain remote model agents preferably by having the model agents pushed to the domain, although other forms of subscription are possible. The model agents broadly enable inter-domain and inter-enterprise analysis and optimization and can embody key features of permissibility (only allowed portions are sent), pushed subscription (each domain pushes changes if they occur), and hybrid composition (both state and behavior are transferred).

30

Inter-domain connectivity plane

FIGURE 10A is a block diagram of many-to-many

interaction between domains, and FIGURE 10B is a block diagram of simplified domain interaction enabled by the inter-domain connectivity plane according to the present invention. As shown in FIGURE 10A, an environment 250 can include a first domain 252 and a second domain 254. Domain 252 can include a number of applications 256, 258 and 260. Similarly, domain 254 can include a number of applications 262 and 264. In the many-to-many scheme of environment 250, applications 256, 258 and 260 of domain 252 must handle data conversions between applications 262 and 264 of domain 254, and vice versa. Further, each application has concerns about security and permissibility, and multiple inter-domain communication channels must be maintained.

FIGURE 10B is a block diagram of simplified domain interaction enabled by the inter-domain connectivity plane of the present invention. As shown, an environment 270 can include a first domain 272 and a second domain 274. Domain 272 can include a number of applications 276, 278 and 280, and domain 274 can include applications 282 and 284. As has been described, data and information from domains 272 and 274 can be abstracted and provided to engines 286 and 288. Engines 286 and 288 then interact on the inter-domain connectivity plane thus simplifying the interaction between domain 272 and 274. The inter-domain connectivity plane, upon which various domain engines 286 and 288 can interact, thereby allows multiple, diverse entities (e.g., enterprises, divisions, etc.) to link supply chain operations.

FIGURE 11 is a block diagram of one embodiment of an inter-domain connectivity plane according to the present invention. As shown, an environment 290 includes a native plane 292. Native plane 292 can include native sources 294

associated with a first domain and native sources 296 associated with a second domain. Native sources 294 and 296 can comprise data sources, planning engines and other domain related data and applications. Data and information from native sources 294 and 296 are abstracted onto an inter-domain connectivity plane 298 using various adapters within the universal adapter framework as has been described. For example, business object servers (BOS's) can be used to raise data to inter-domain connectivity plane 298. A domain engine 300 can be associated with the first domain and receives data and information abstracted from native sources 294. In an analogous manner, a domain engine 302 can be associated with the second domain and receives data and information abstracted from native sources 296. Inter-domain connectivity plane 298 can further include one or more domain engines 304 not specifically associated with a domain but operating on inter-domain connectivity plane 298 to provide or perform desired functions. Inter-domain connectivity plane 298 enables global messaging, as represented by line 306, directly between domain engines 300, 302 and 304, as viewed by domain engines 300, 302 and 304. However, the messaging is actually transparently supported by native messaging 308 on native plane 292. Inter-domain connectivity plane 298 provides an abstracted layer to allow harmonization across distributed domains and provides significant advantages by harmonizing such aspects as object structure, messaging paradigm, naming and security.

FIGURE 12 is a block diagram of a hub and spoke collaboration network formed by engines on the inter-domain connectivity plane. As shown, an inter-domain connectivity plane 310 can include a plurality of domain engines 312 and

314. In the hub and spoke arrangement, engines 312 and 314 are associated with one of two kinds of domains or entities in the network. As shown, there are hub engines 312 and spoke engines 314. Hub engines 312 can communicate 5 simultaneously with multiple other hub engines 312. Spoke engines 314, on the other hand, only communicate with parent hub engines 312. Additionally, in one implementation, spoke engines 314 do not perform analysis but can only feed information to and be fed information from their parent hub engine 312. In another 10 implementation, spoke engines 314 can be mid tier or web interface tier engines. In this scheme, the mid tier engines can perform some lightweight analysis, and the web interface tier engines do not perform analysis. Thus, the 15 arrangement would enable three forms of communication within inter-domain connectivity plane 310: hub-to-hub, hub-to-spoke and user interface-to-hub.

The inter-domain connectivity plane enables inter-domain visibility, signaling, collaboration, analysis and 20 optimization. Inter-domain visibility is a basic function of the inter-domain connectivity plane and gives visibility into data of other domains. Data can be pushed as well as pulled and can be used for user interaction as well as automatic analysis purposes. Inter-domain signaling allows 25 domains to signal other domains. For example, the inter-domain connectivity plane can enable an early warning system in which a domain can signal other domains when certain conditions arise. The inter-domain connectivity plane also facilitates engine-to-engine collaboration 30 across domains (as opposed to simple person-to-person collaboration). Further, the inter-domain connectivity plane can enable inter-domain analysis and optimization in

part by enabling the transfer of the model agents described above.

The inter-domain connectivity plane is built upon the core environment and, as such, possesses the attributes and components of the core environment. Of these components, global messaging (used to message between multiple domains) and transactional messaging (used to keep the domains transactionally consistent) is important. Business object servers used to raise objects from the native plane to the inter-domain connectivity plane are also important. The engines on the inter-domain connectivity plane provide the bulk of the function and support numerous functions, including the early warning system, multi-domain collaboration, inter-domain analysis and optimization, a permissibility framework, and global naming. The engines support the model agents which enable partial mirroring of remote models and support publish and subscribe model agent distribution. Security is also a fundamental attribute of the inter-domain connectivity plane. At the lowest level, the inter-domain connectivity plane inherits the security features of the core environment including: user interface to engine security (authentication and encryption) and messaging security (authentication and encryption). At the design level, security is ensured by a combination of permissibility and push technology.

Scalability is an important aspect, and the inter-domain layer is scaleable in several dimensions. Guaranteed asynchronous messaging increases scalability by minimizing network dependencies. This can be especially critical as the number of communicating domains increases. Scalability is also achieved by avoiding the many-to-many problem at the application level (security, permissibility

and data conversion). Scaleability is provided by allowing a single domain to talk to multiple other domains with no controlling domain. A domain typically needs to communicate with multiple other domains. However, the inter-domain connectivity plane allows the domains to do so without the need for any sort of higher level controller. Additionally if communications between a pair of domains breaks down, communications can continue between other domains, leading to far greater scaleability.

10

#### Global naming

The inter-domain connectivity plane can implement a global naming scheme that allows entities to be uniquely defined on a global basis. For example, the global naming scheme can comprise a three part naming scheme using the Internet domain name server (DNS) at the highest level, enterprise name services (such as CORBA or LDAP) at the next level, and transient name services at the process level. With respect to naming, both name wrapping and name mapping provide harmonization. The name wrapping allows non-global names to be raised into a global name space and ensures unique naming when data is brought up to the inter-domain connectivity plane. Name mapping, on the other hand, provides a map to ensure that different names for the same object refer to the same object. Name mapping can allow two global names to refer to the same entity by mapping the names to each other.

30

Although the present invention has been described in detail, it should be understood that various changes, substitutions and alterations can be made hereto without departing from the spirit and scope of the invention as defined by the appended claims.

WHAT IS CLAIMED IS:

1. A computer implemented system for inter-domain interaction, comprising:

a native plane comprising:

5 a first native source associated with a first domain; and

a second native source associated with a second domain;

10 the first and second native sources including domain related data and applications;

a shared inter-domain connectivity plane comprising:

a first domain engine associated with the first domain; and

15 a second domain engine associated with the second domain;

the first and second domain engines collaborating to perform domain analysis functions; and

20 adapters associated with the first and second native sources, the adapters operating to abstract data and behavior from the first and second native sources onto the shared inter-domain connectivity plane;

25 the first domain engine receiving data and information abstracted from the first native source via adapters and the second domain engine receiving data and information abstracted from the second native source and adapters; and

the first and second domain engines operating to perform global messaging between one another transparently supported by native messaging on the native plane.

30 2. The system of Claim 1, wherein the adapters operate within a universal adaptor framework to abstract data and behavior onto the inter-domain connectivity plane.

3. The system of Claim 2, wherein the adapters include a business object server that operates to raise data and functionality onto the inter-domain connectivity plane.

4. The system of Claim 1, wherein the inter-domain connectivity plane further comprises a third domain engine not specifically associated with a native source or domain, the third domain engine operating on the inter-domain connectivity plane to perform inter-domain analysis functions and data distribution.

5. The system of Claim 1, further comprising a plurality of additional domain engines associated with additional domains, the domain engines forming a hub and spoke collaboration network.

6. The system of Claim 5, wherein the domain engines can be hub engines or spoke engines, the hub engines can communicate simultaneously with multiple other hub engines, and the spoke engines only can communicate with a parent hub engines.

7. The system of Claim 6, wherein the spoke engines feed information to and can be fed information from the parent hub engine and do not perform analysis themselves.

8. The system of Claim 1, further comprising:  
additional native sources associated with the first  
domain; and

5 additional native sources associated with the second  
domain;

the additional native sources including domain related  
data and functionality; and

10 the adapters associated with the additional native  
sources and operating to abstract data and information from  
the additional native sources onto the inter-domain  
connectivity plane.

9. The system of Claim 1, where in the first and  
second domains comprise separate enterprises.

15

10. The system of Claim 1, where in the first and  
second domains comprise separate supply chain entities  
within a single enterprise.

11. A computer implemented process for inter-domain interaction, comprising:

5 establishing a native plain comprising a plurality of native sources respectively associated with domains, the plurality of native sources including domain related data and applications;

10 establishing an inter-domain connectivity plane comprising a plurality of domain engines respectively associated with the domains, the plurality of domain engines operating to perform domain analysis functions;

abstracting data and information from the plurality of native sources onto the inter-domain connectivity plane;

15 receiving, at the plurality of domain engines, data and functionality abstracted from the plurality of native sources; and

performing messaging between the plurality of domain engines transparently supported by native messaging on the native plane.

20 12. The process of Claim 11, wherein abstracting is accomplished by adapters operating within a universal adaptor framework to abstract data and information onto the inter-domain connectivity plane.

25 13. The process of Claim 11, wherein abstracting is accomplished, in part, by business object servers that operate to raise data onto the inter-domain connectivity plane.

14. The process of Claim 11, wherein establishing the inter-domain connectivity plane further comprises establishing a domain engine not specifically associated with a native source or domain which operates on the inter-domain connectivity plane to perform inter-domain analysis functions.

15. The process of Claim 11, wherein establishing the inter-domain connectivity plane comprises arranging the plurality of domain engines to form a hub and spoke collaboration network.

16. The process of Claim 15, wherein the domain engines can be hub engines or spoke engines, the hub engines can communicate simultaneously with multiple other hub engines, and the spoke engines only can communicate with a parent hub engines.

17. The process of Claim 16, wherein the spoke engines comprise either mid tier engines or web interface tier engines, the mid tier engines operable to perform some analysis and the web interface tier engines not operable to perform analysis.

18. The process of Claim 16, wherein the spoke engines feed information to and can be fed information from the parent hub engine and do not perform analysis themselves.

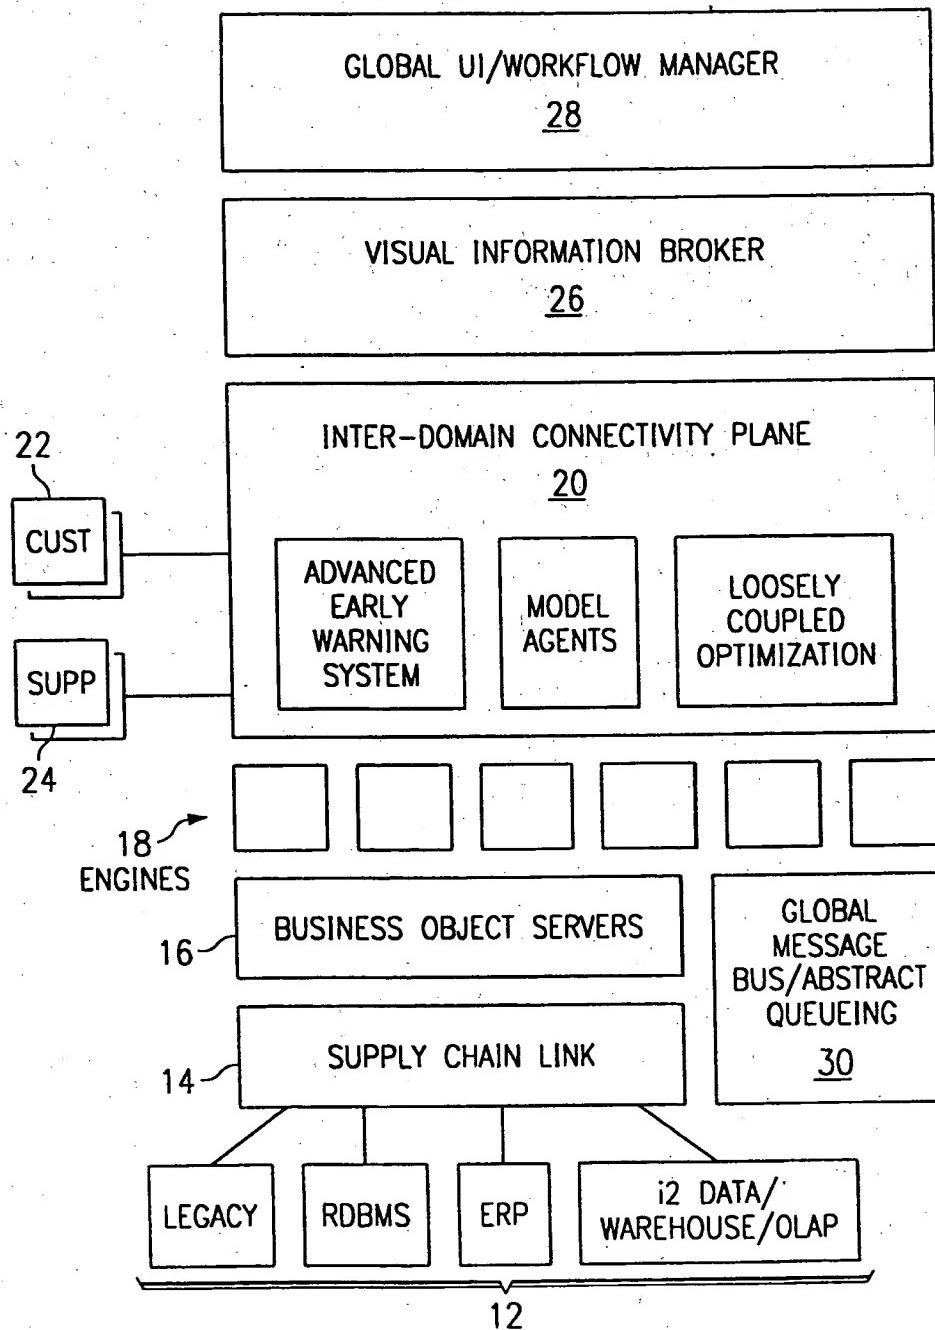
19. The process of Claim 11, where in a portion of the domains comprise a plurality of separate manufacturing enterprises.

20. The process of Claim 11, where in a portion of the domains comprise a plurality of manufacturing entities within a single enterprise.

1/7

10

FIG. 1



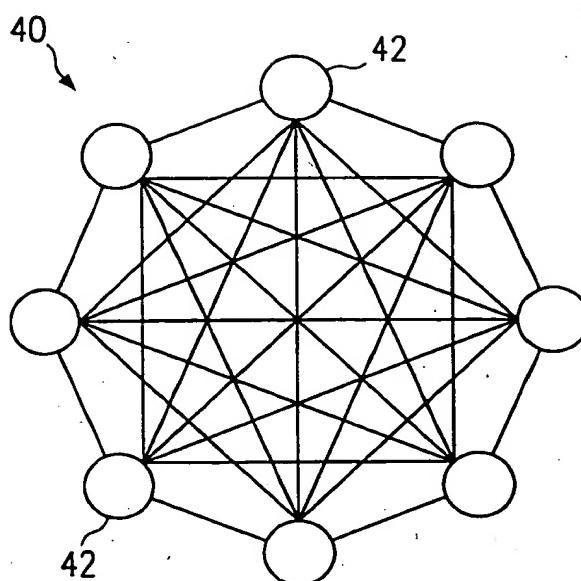


FIG. 2A

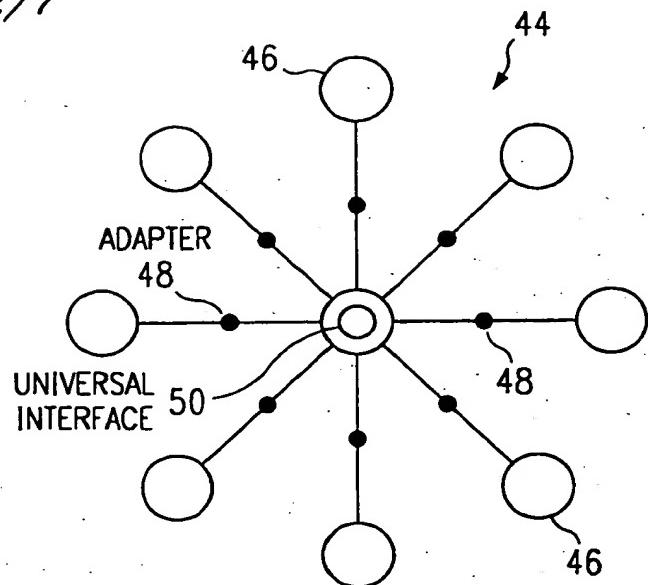


FIG. 2B

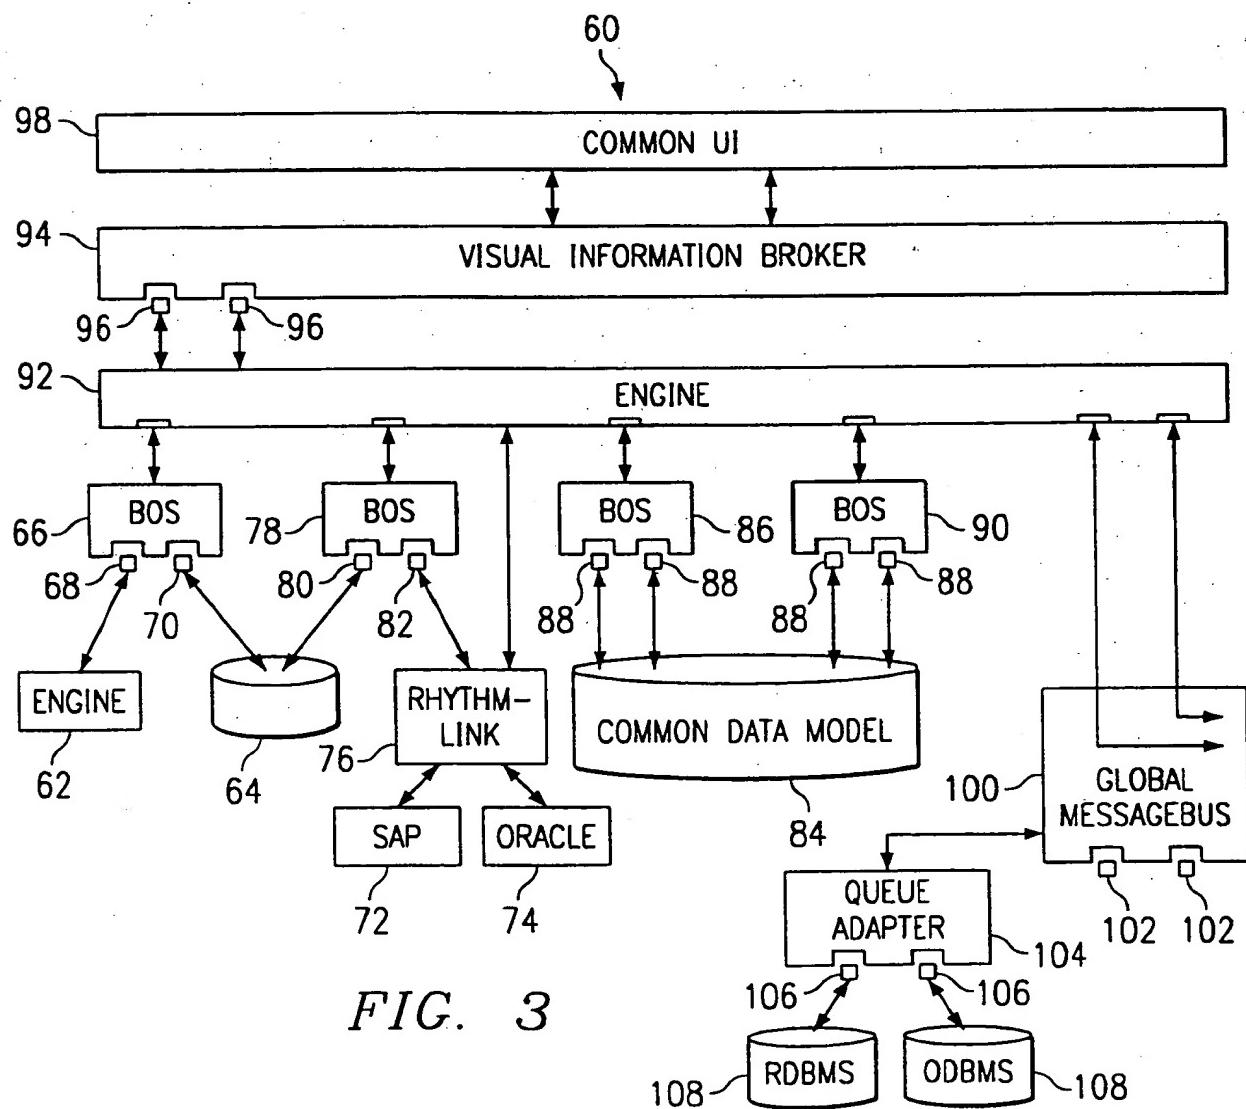


FIG. 3

3/7

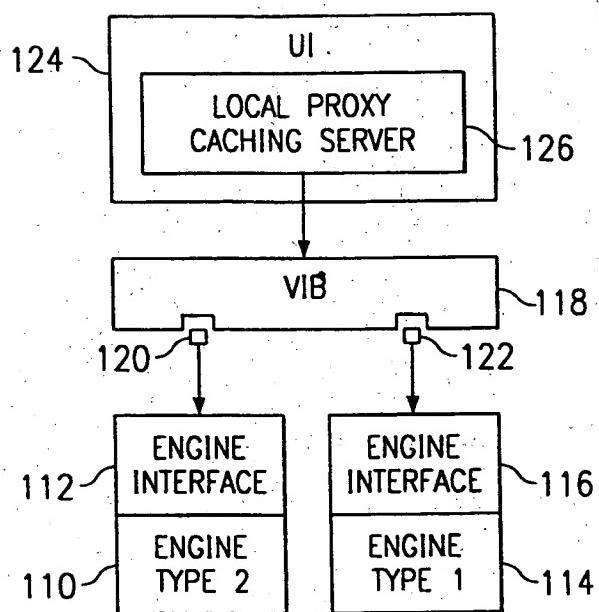


FIG. 4

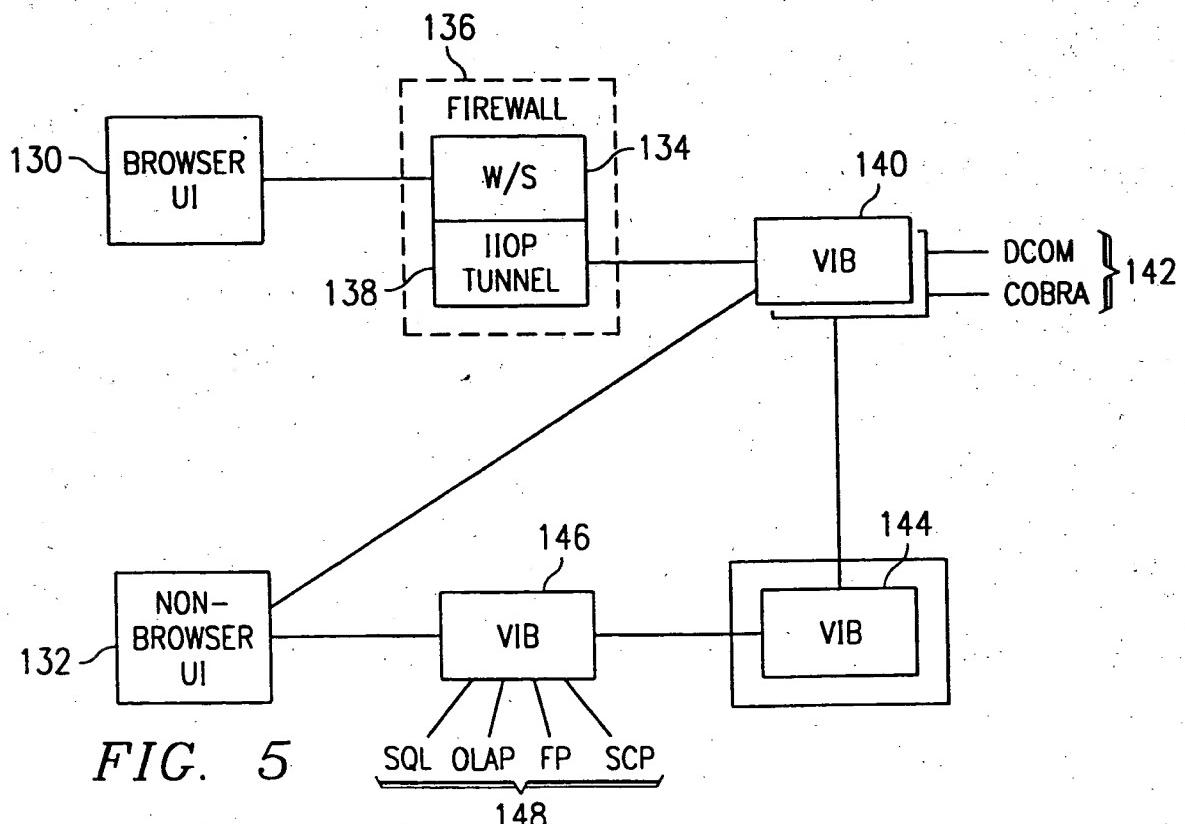


FIG. 5

4/7

FIG. 6

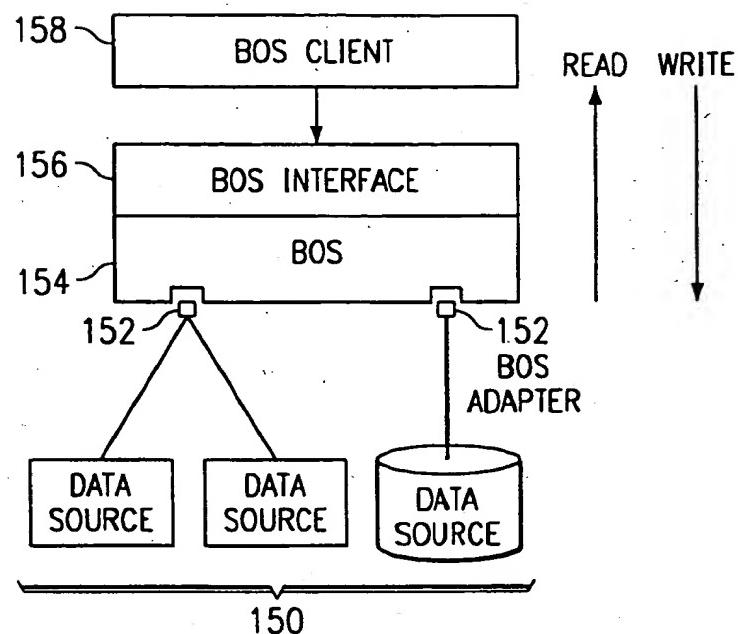
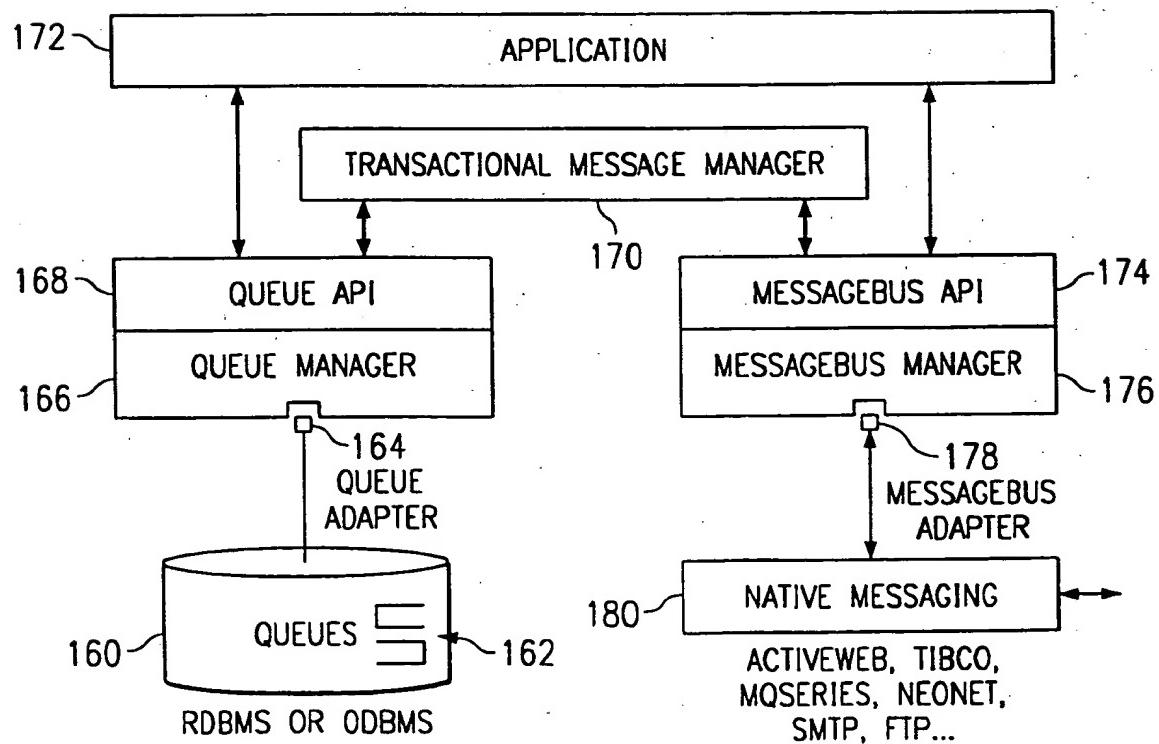
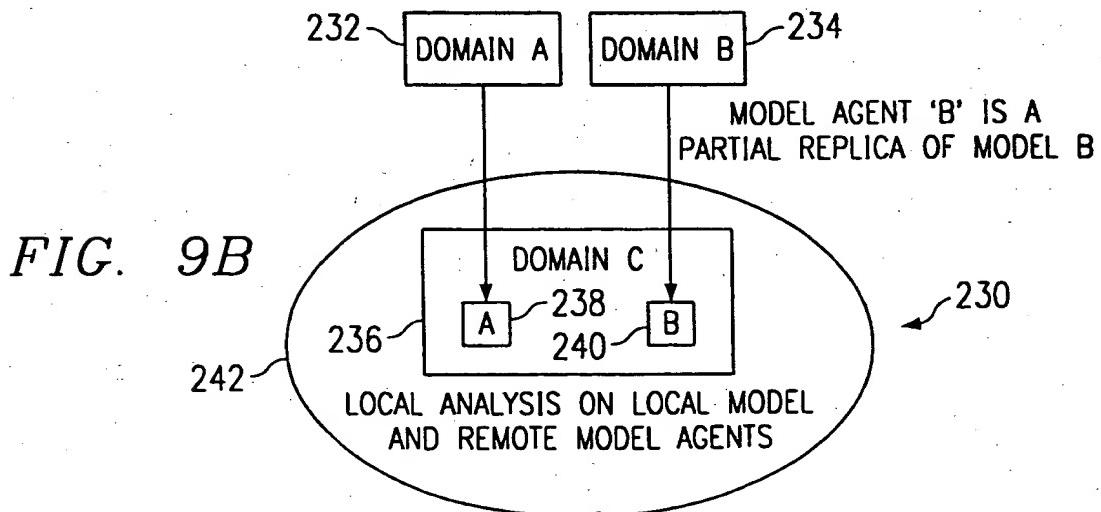
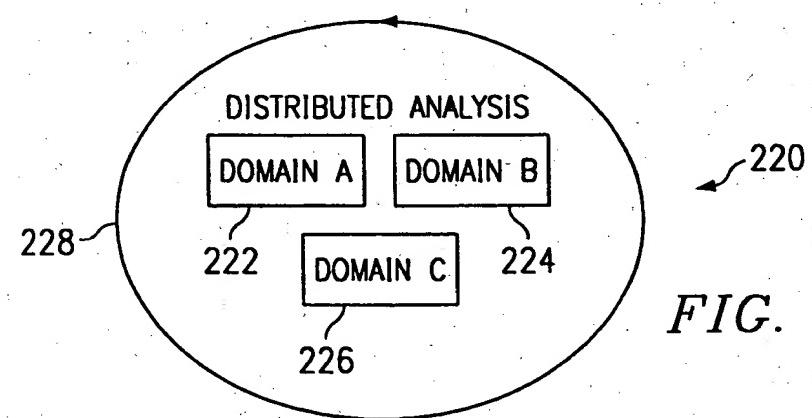
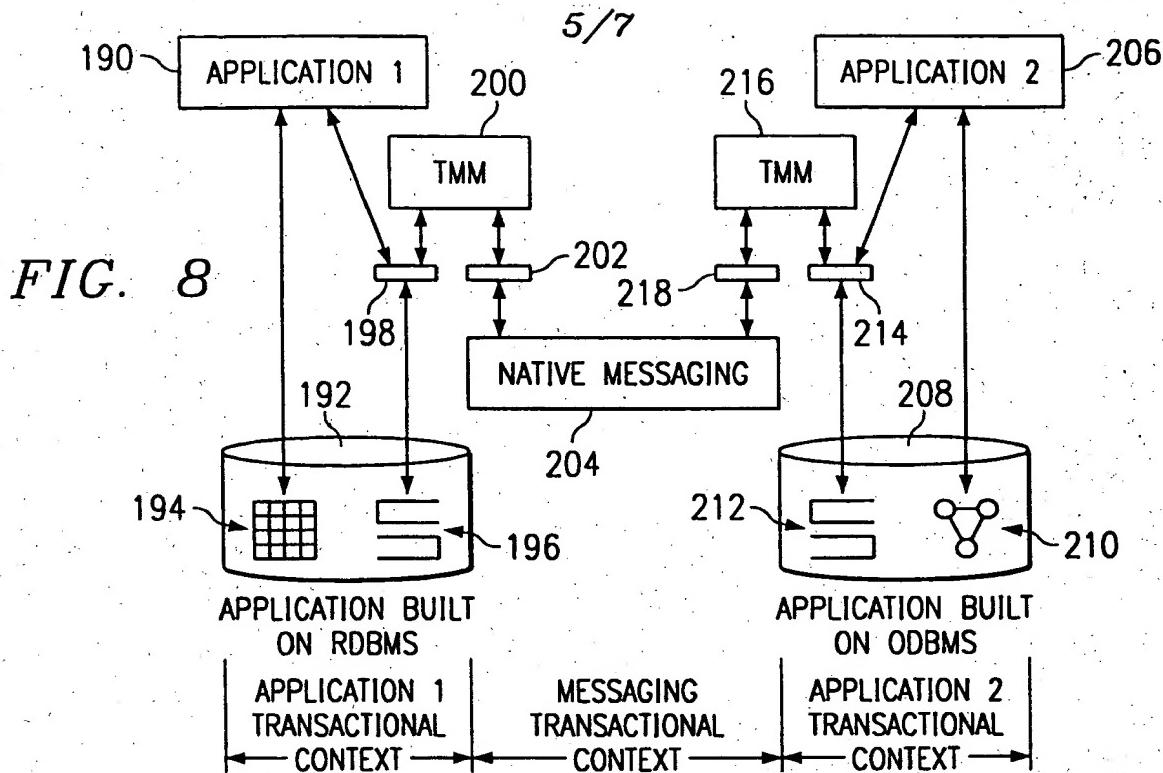


FIG. 7





6/7

FIG. 10A

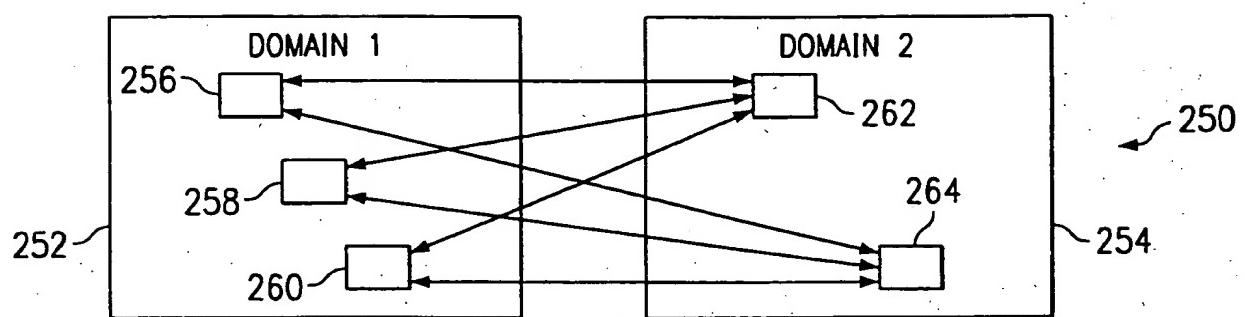
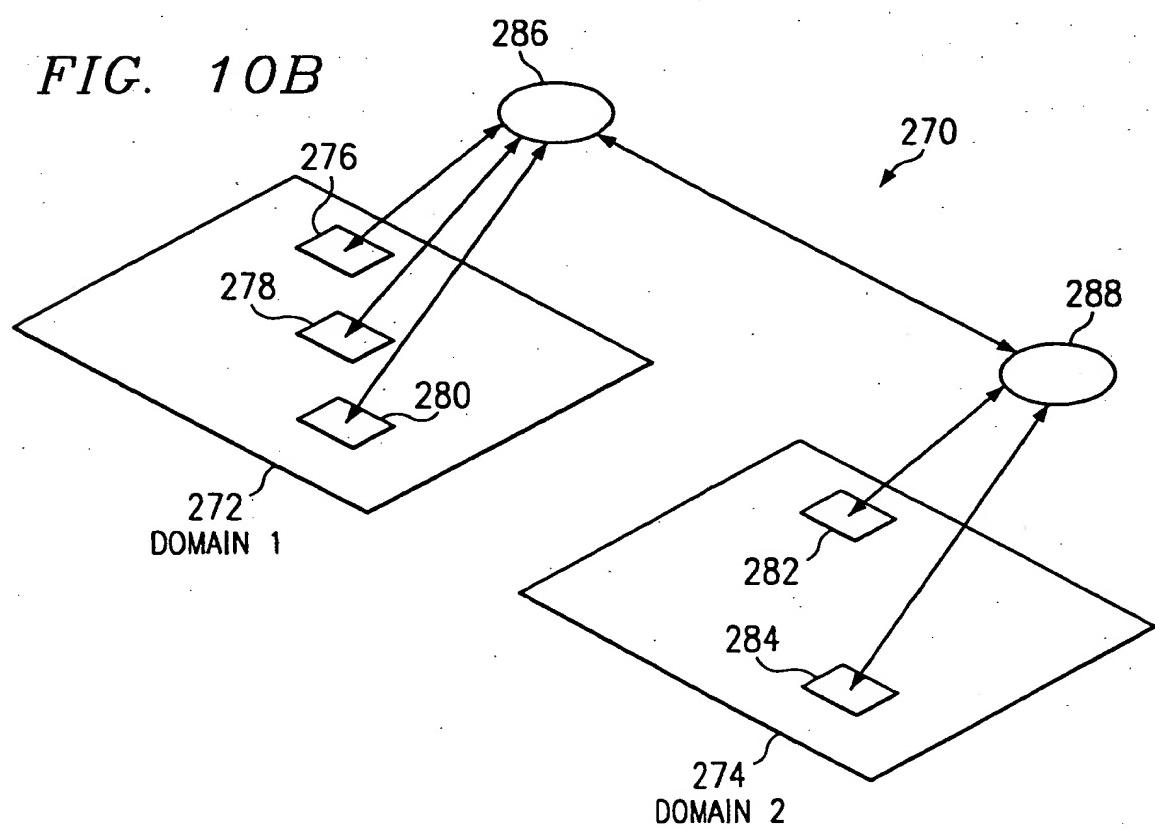


FIG. 10B



7/7

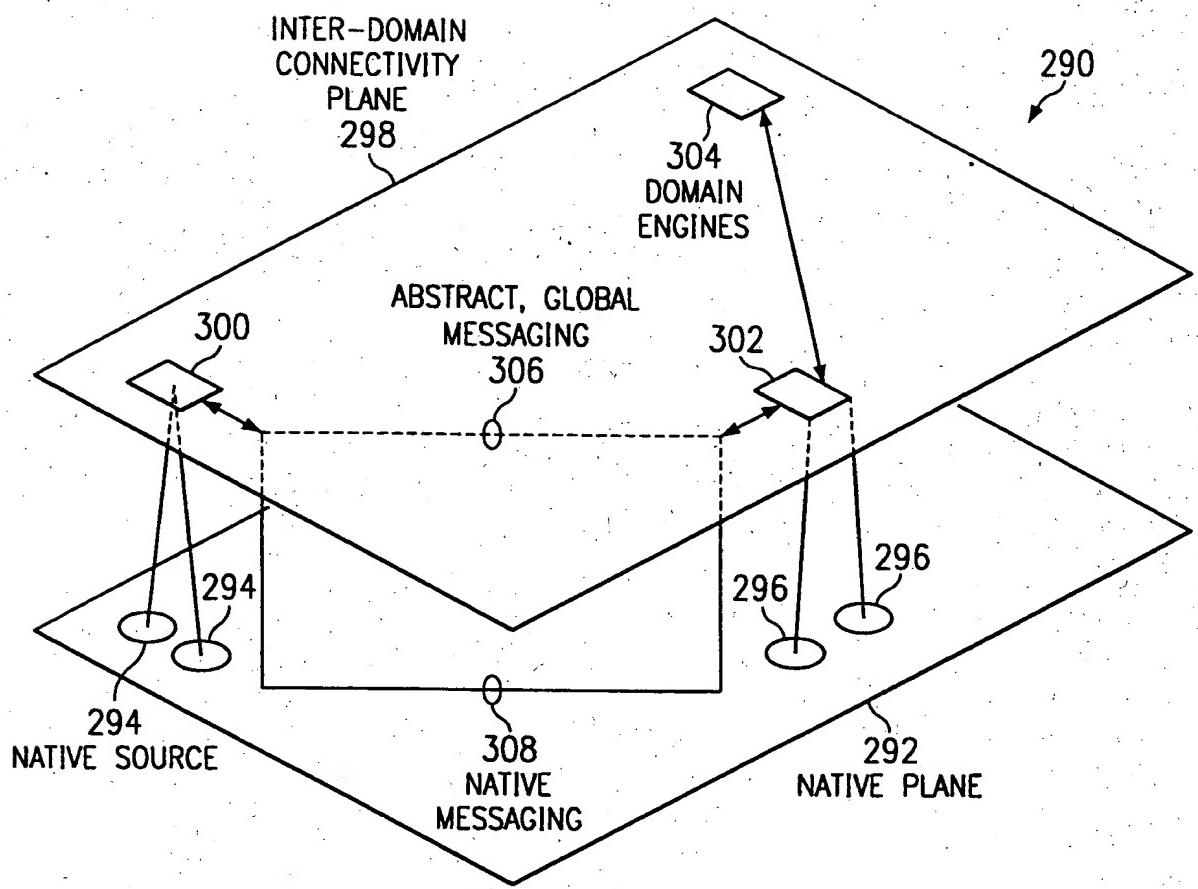


FIG. 11

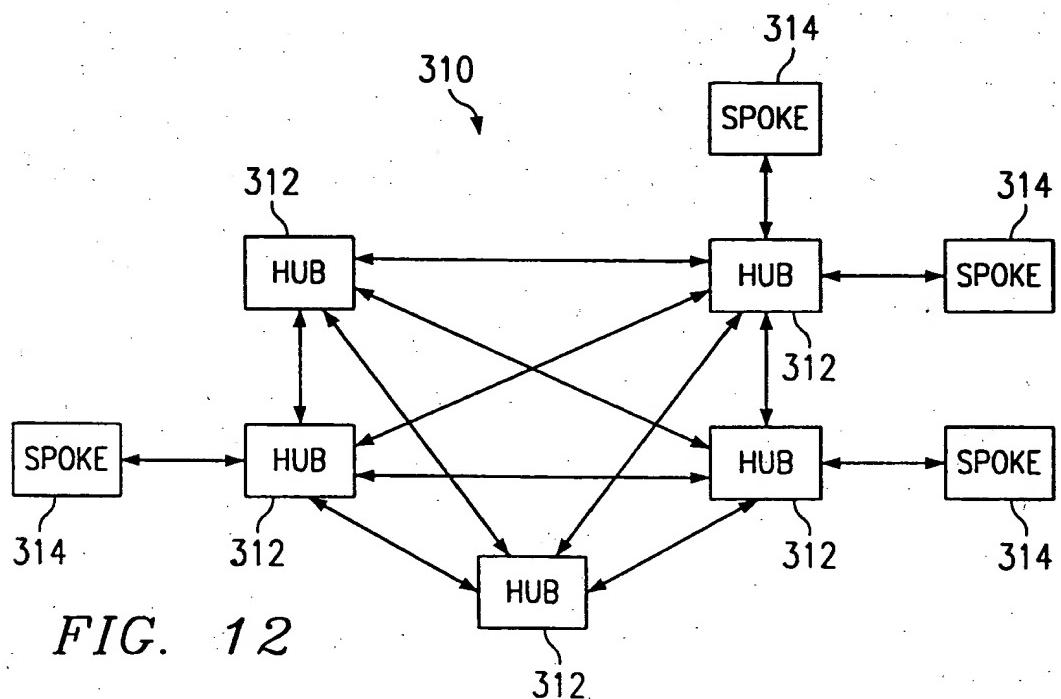


FIG. 12

# INTERNATIONAL SEARCH REPORT

Internat. application No

PCT/US 98/17359

**A. CLASSIFICATION OF SUBJECT MATTER**  
IPC 6 G06F17/60

According to International Patent Classification (IPC) or to both national classification and IPC

**B. FIELDS SEARCHED**

Minimum documentation searched (classification system followed by classification symbols)

IPC 6 G06F

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

**C. DOCUMENTS CONSIDERED TO BE RELEVANT**

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	<p>BARBUCEANU M ET AL: "Agent based design and simulation of supply chain systems"          PROCEEDINGS SIXTH IEEE WORKSHOPS ON ENABLING TECHNOLOGIES: INFRASTRUCTURE FOR COLLABORATIVE ENTERPRISES (CAT. NO.97TB100178), PROCEEDINGS OF IEEE 6TH WORKSHOP ON ENABLING TECHNOLOGIES: INFRASTRUCTURE FOR COLLABORATIVE ENTERPRISES, CAMBRIDGE, MA, USA., pages 36-41, XP002090169          ISBN 0-8186-7967-0, 1997, Los Alamitos, CA, USA, IEEE Comput. Soc, USA          see page 36, column 1, line 15 - column 2, line 34          see page 38, column 1, line 2 - page 40, column 1, line 37</p> <p style="text-align: center;">-/-</p>	1-20

Further documents are listed in the continuation of box C.

Patent family members are listed in annex.

\* Special categories of cited documents :

- "A" document defining the general state of the art which is not considered to be of particular relevance
- "E" earlier document but published on or after the international filing date
- "L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
- "O" document referring to an oral disclosure, use, exhibition or other means
- "P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.

"&" document member of the same patent family

Date of the actual completion of the international search

Date of mailing of the international search report

15 January 1999

29/01/1999

Name and mailing address of the ISA  
 European Patent Office, P.B. 5818 Patentlaan 2  
 NL - 2280 HV Rijswijk  
 Tel. (+31-70) 340-2040, Tx. 31 651 epo nl,  
 Fax: (+31-70) 340-3016

Authorized officer

Pedersen, N

# INTERNATIONAL SEARCH REPORT

International Application No.  
PCT/US 98/17359

## C.(Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT

Category	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	<p>BARBUCEANU M ET AL: "Coordinating multiple agents in the supply chain"          PROCEEDINGS OF THE 5TH WORKSHOPS ON ENABLING TECHNOLOGIES: INFRASTRUCTURE FOR COLLABORATIVE ENTERPRISES (WET ICE '96) (96TB100058), PROCEEDINGS OF WET ICE '96. IEEE 5TH WORKSHOP ON ENABLING TECHNOLOGIES; INFRASTRUCTURE FOR COLLABORATIVE ENTERPRISES, ST, pages 134-141, XP002090170 ISBN 0-8186-7446-6, 1996, Los Alamitos, CA, USA, IEEE Comput. Soc. Press, USA see page 135, column 1, line 3 - line 46          see figure 1</p> <p>---</p>	1-20
X	<p>FOX, M.S., CHIONGLO, J.F., AND BARBUCEANU, M.: "The Integrated Supply Chain Management System"          7 December 1993, DEPT. OF INDUSTRIAL ENGINEERING, UNIVERSITY OF TORONTO.          XP002090173          Available on the internet on the 13th January 1999  <a href="http://128.100.52.15/EIL/iscm-descr.html">http://128.100.52.15/EIL/iscm-descr.html</a>          see page 4, line 16 - page 9, line 3</p> <p>---</p>	1-20
P,X	<p>FU-REN LIN ET AL: "Modeling supply-chain networks by a multi-agent system"          PROCEEDINGS OF THE THIRTY-FIRST HAWAII INTERNATIONAL CONFERENCE ON SYSTEM SCIENCES (CAT. NO.98TB100216), PROCEEDINGS OF THE THIRTY-FIRST HAWAII INTERNATIONAL CONFERENCE ON SYSTEM SCIENCES, KOHALA COAST, HI, USA, 6-9 JAN. 1998, pages 105-114 vol.5, XP002090171          ISBN 0-8186-8255-8, 1998, Los Alamitos, CA, USA, IEEE Comput. Soc, USA          see page 107, column 2, line 24 - page 110, column 2, line 29</p> <p>---</p>	1-20
A	<p>EP 0 770 967 A (PHILIPS ELECTRONICS NV)          2 May 1997          see page 3, line 1 - page 4, line 8</p> <p>---</p>	1-20
A	<p>HIRSCH B E ET AL: "Information system concept for the management of distributed production"          COMPUTERS IN INDUSTRY, vol. 26, no. 3, August 1995, page 229-241          XP004002562          see page 233, column 2, line 17 - page 235, column 2, line 16          see figure 1</p> <p>---</p>	1-20

-/-

## INTERNATIONAL SEARCH REPORT

International	Application No
PCT/US 98/17359	

## C.(Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	<p>CIANCARINI P ET AL: "A multi-agent process centered environment integrated with the WWW"          PROCEEDINGS SIXTH IEEE WORKSHOPS ON ENABLING TECHNOLOGIES: INFRASTRUCTURE FOR COLLABORATIVE ENTERPRISES (CAT. NO.97TB100178), PROCEEDINGS OF IEEE 6TH WORKSHOP ON ENABLING TECHNOLOGIES: INFRASTRUCTURE FOR COLLABORATIVE ENTERPRISES, CAMBRIDGE, MA, USA,, pages 113-120, XP002090172 ISBN 0-8186-7967-0, 1997, Los Alamitos, CA, USA, IEEE Comput. Soc, USA          see page 1, column 1, line 1 - column 2, line 27</p> <p>---</p> <p>EP 0 615 198 A (IBM) 14 September 1994          see column 1, line 1 - column 2, line 10</p> <p>---</p>	1-20
A		1-20

# INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No

PCT/US 98/17359

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
EP 0770967 A	02-05-1997	NONE	
EP 0615198 A	14-09-1994	NONE	